TITLE: Naturally occurring combinations of receptors from single cell transcriptomics in endothelial cells

Sergii Domanskyi[1], Alex Hakansson[2], Michelle Meng[2], Joshua S. Graff Zivin[3], Carlo Piermarocchi[1], Giovanni Paternostro[2], Napoleone Ferrara[4]

[1]Department of Physics and Astronomy, Michigan State University, East Lansing, MI, 48824, USA
[2]Sanford Burnham Prebys Medical Discovery Institute, La Jolla, CA, 92037, USA
[3]School of Global Policy and Strategy and Department of Economics, University of California, San Diego 9500 Gilman Drive, MC 0519, La Jolla, CA 92093, USA
[4]Moores Cancer Center, University of California San Diego, La Jolla, CA 92093, USA

Corresponding authors:
domansk6@msu.edu
jgraffzivin@ucsd.edu
piermaro@msu.edu
giovanni@sbpdiscovery.org
nferrara@health.ucsd.edu

ABSTRACT

VEGF inhibitor drugs have been successful, especially in ophthalmology, but not all patients respond to them. Combinations of drugs are likely to be needed for a really effective therapy of angiogenesis-related diseases. In this paper we describe naturally occurring combinations of receptors in endothelial cells that might help to identify targets for drug combinations. We also develop and share a new computational method and a software tool called DECNEO to identify them.

Single-cell gene expression data are used to identify a set of co-expressed endothelial cell receptors, conserved among species (mice and human) and enriched, within a network, of connections to up-regulated genes. This set includes several receptors previously shown to play a role in angiogenesis. The reproducibility, evolutionary conservation and role in angiogenesis of the naturally occurring combinations of receptors are supported by multiple highly significant statistical tests from large datasets, including an independent validation set. We also show tissue-specific combinations and, in the case of choroid endothelial cells, consistency with both well-established and recent experimental findings, presented in a separate experimental paper.

The results and methods presented here allow one small step forward in the understanding of the syntax of intercellular communication.

AUTHOR SUMMARY

Intercellular communication is essential for animal development and function and therefore it must have already been present when the first animals evolved, more than 600 million years ago. Many effective drugs act by modifying messages between cells.

Thus, a fuller understanding of the language of cells should lead to the design of a wider range of safer, more effective drugs and drug combinations.

In this paper we show an example of remarkable similarity in patterns of cellular communication between mice and humans, two species that began to diverge from their common ancestor approximately 75 million years ago. This paper focuses on combinations of signals important for the growth of blood vessels. The methods used, however, are of a general nature and take advantage of recent technology that is able to measure the quantities of molecules within a single cell. We hope that colleagues with expertise in different fields, science students, and people contemplating a career in science will be inspired by the prospect of one day fully understanding the ancient language of cells.

INTRODUCTION

Angiogenesis is the development of new blood vessels. It is a physiological process required for embryonic development, adult vascular homeostasis and tissue repair (Chung and Ferrara, 2011). It is a fundamental process of human biology involved in a large number of diseases, including cancer, eye diseases, arthritis, psoriasis and many other ischemic, inflammatory, infectious, and immune disorders.

During tumor progression, the new vessels provide neoplastic tissues with nutrients and oxygen; in intraocular vascular disorders, growth of abnormal, leaky blood vessels may destroy the retina and lead to blindness (Chung and Ferrara, 2011; Potente et al., 2011). Extensive efforts to dissect the molecular basis of angiogenesis and to identify therapeutic targets resulted in the discovery of key signaling pathways involved in vascular development and differentiation (Chung and Ferrara, 2011; Yancopoulos et al., 2000). In particular, numerous studies have established the pivotal role of the VEGF (Vascular Endothelial Growth Factor) pathway in physiological and pathological angiogenesis and therapies targeting this pathway have achieved success in treatment of cancer and ocular disorders (Apte et al., 2019; Ferrara and Adamis, 2016). VEGF inhibitors are currently the main class of FDA approved anti-angiogenic drugs and they are used as a therapy for cancer and for eye diseases leading to blindness. One of us has provided fundamental contributions to the understanding of the role of VEGF in angiogenesis and to the development of these drugs (Ferrara and Adamis, 2016).

VEGF inhibitor drugs have been successful, especially in ophthalmology, but not all patients respond to them (Apte et al., 2019; Chung and Ferrara, 2011; Ferrara and Adamis, 2016). Multiple extracellular factors are known to control of angiogenesis and combinations of drugs acting on these molecules are likely to be needed for a really effective therapy of angiogenesis-related diseases (Apte et al., 2019; Chung and Ferrara, 2011; Ferrara and Adamis, 2016).

Several methods are currently used to discover novel drug combinations (Feala et al., 2010). Beyond empirical combinations of individually effective drugs, they include the use of drug screens, using brute force (that is, for small sets, exhaustive testing of all possible combinations) or search algorithms depending on the phenotypic response (Kang et al., 2014). They also include modeling of biological systems, based on statistical associations or on direct modeling of molecular interactions (Ong et al., 2015). These

approaches can also be combined, but our knowledge of the details of *in vivo* biological systems is still incomplete and progress to the clinic is therefore slow and inefficient. We also know that gene products are co-expressed in sets that have specific functions and these might represent potential drug targets for drug combinations (Feala et al., 2012). Understanding the regulation and role of these co-expressed sets has been a key focus of biological investigations, starting from the discovery of the *lac* operon in bacteria by Jacobs and Monod (Jacob et al., 1960).

In this paper we describe naturally occurring combinations of genes or their products, for which we use the abbreviation comberon, by analogy with the operon.

The increasing number of single-cell RNA-seq datasets that provide information on organ and tissue-specific signaling and dedicated initiatives such as the Human Cell Atlas, the NIH Human Biomolecular Atlas Program and the EU Lifetime Initiative (HuBmap, 2019; Regev et al., 2017) aiming to map all the cell types of the human body, offer a new source of data that can inform the understanding of co-expression at the single cell level. Co-expression at the single cell level gives a more detailed understanding of biological regulation compared to bulk measurements.

Evolution by natural selection is the main organizing principle in biology (Nurse, 2020). The evolutionary age and rate of change of genes can provide information about their functions and interactions (Szedlak et al., 2016). We expect that any functionally important co-expression of genes would be conserved during evolution. Considerations analogous to these have guided combinatorial interventions before, for example the finding by Shin Yamanaka and his collaborators that delivering the same four transcription factors can reprogram cells to a pluripotent state in both mice and humans (Karagiannis et al., 2019).

Furthermore, while incomplete, network models of intracellular interactions are being developed using an increasing number of datasets and can contribute to the validation of new findings (Ong et al., 2015). Endothelial cells, the main targets of VEGF, are the key cell type in angiogenesis and their receptors integrate multiple signals to decide if angiogenesis should be activated (Ferrara et al., 2003; Ferrara and Kerbel, 2005). We show how to identify receptors linked, within network models, to significantly up-regulated genes by calculating a measure called network enrichment. It is a measure of statistical association, but, in some cases, it might represent downstream changes due to the signals transmitted by the receptors.

We show here that an integrated analysis of multiple single cell RNA-seq datasets, using co-expression, level of expression, evolutionary conservation and network enrichment can be used to identify naturally occurring combinations (comberons) of receptors in endothelial cells. This set includes VEGF receptors and other receptors previously shown to play a role in angiogenesis. Figure 2 provides a visual introduction to our analysis and to key results.

We then show that comberons are in part tissue-specific and that, in the case of choroid endothelial cells, results are consistent with both well-established and recent experimental findings, presented in a separate experimental paper (Li et al., 2021).

We also introduce a new software tool, DECNEO (DEndrogram of Co-expression with Evolutionary Conservation and Network Enrichment Ordering), that identifies general and tissue-specific naturally occurring combinations (comberons) of receptors in endothelial cells and provides visualization and statistical analysis.

RESULTS

Overview of single cell gene expression data

We used the DCS pipeline (Domanskyi et al., 2021; Domanskyi et al., 2019b) to process the single-cell RNA-seq PanglaoDB database (Franzén et al., 2019). The pipeline was applied to each sample independently, without correcting for batch effects, to normalize the data and identify endothelial and non-endothelial cells. Figure 1A shows 101,470 *Mus musculus* endothelial cells from 397 samples within 51 studies projected on a t-SNE layout color-coded by tissue of origin, following the Panglao tissue labels. The proportion of endothelial cells represented by each tissue type is shown in the bar plots. A large fraction of the cells in this dataset is from brain tissues. Figure 1B summarizes the *Homo sapiens* portion of the PanglaoDB endothelial cells identified by DCS. There are much more data from *Mus musculus* studies than *Homo sapiens* resulting in higher statistical power in the analysis of *Mus musculus* tissues. In the control and validation analysis we also use the PanglaoDB original annotation of endothelial cell by the Alona software (Franzén and Björkegren, 2020). The Alona analysis of the PanglaoDB database results in a different number of samples containing endothelial cells. We will refer to the PanglaoDB data processed using DCS and Alona as "PanglaoDB-DCS" and "PanglaoDB-Alona", respectively.

In the tissue-specific analysis we use the single cells dataset from Voigt et al. (Voigt et al., 2019) containing *Homo sapiens* choroid data from 18,621 cells in 8 samples. The samples were enriched for endothelial cells, mapped by CellRanger and then normalized and annotated by DCS. Figure S10A-B shows a t-SNE layout and the relative proportions of choroid cells colored by cell type. Although this dataset contains 8 samples, there are no strong batch effects, and the main group of cells, primarily endothelial, includes all 8 biological samples, Figure S10C.

Figure 1. 123,260 endothelial cells identified with DCS from PanglaoDB, which collects 22,382,464 cells in total from 1,368 samples, in 236 studies. The panels show the t-SNE projection and bar chart of the relative fractions of the endothelial cells, color-coded by tissue of origin, separately for: (A) *Mus musculus*, and (B) *Homo sapiens*. The asterisk after the tissue names indicates that the PanglaoDB data include multiple subtypes of tissues combined into one. The DCS processing pipeline was used for each of the 1,368 samples independently.

Dendrogram of co-expression ordering

In Figure 2 we show the dendrogram obtained from the analysis of the PanglaoDB-DCS dataset by agglomerative clustering of the receptor co-expression. We also show

the three measures introduced above and their combination. The receptor ordering groups co-expressed receptors close to each other.

The "Known angiogenesis receptors" that are expressed in endothelial cells in the PanglaoDB-DCS *Mus musculus* datasets are shown in a panel below the dendrogram in Figure 2A. These known angiogenesis regulators localize in the ordered dendrogram near each other and near KDR and FLT1, the two main VEGF receptors. The Likelihood-Ratio test p-value of a logistic regression for the distance of the known angiogenesis receptors from KDR and FLT1 is less than $7 \cdot 10^{-4}$. The dendrogram ordering therefore indicates the role of these receptors in angiogenesis. The non-parametric robust PPMLE was also applied to test the distance of the known angiogenesis receptors from KDR and FLT1 in the dendrogram. Both tests gave a p-value $< 5.4 \cdot 10^{-6}$.

Figure 2. The top of panels A and B displays the dendrogram ordering by single cell co-expression of receptors in endothelial cells for: (A) *Mus musculus,* and (B) *Homo sapiens*. Receptors that are co-expressed in individual cells are closer to each other in the dendrogram order. Known angiogenesis receptors are labelled under the dendrogram. The red pointer indicates KDR, the main VEGFA receptor.

The panels immediately below the dendrogram show 3 measures: Expression fraction, Network enrichment and Evolutionary conservation. The lowest panel shows the combination of these 3 measures. The orange solid line in each panel is a moving window average. The largest peak in the combination of measures panel, representing the main comberon (abbreviation for naturally occurring combination) is shown with a red background.

Genes in the main peaks (comberons) and bootstrap analysis

Using the "Combination of measures" of the PanglaoDB-DCS *Mus musculus* data we determined which of the 898 expressed receptors are localized in the largest peak identified by this quantity. Receptors that belong to the largest peak are highlighted in blue in Figure S4 and are listed there. These 38 genes include 10 of the 22 known angiogenesis regulators, among them KDR and FLT1, the two main VEGF receptors. The probability of such an outcome by chance is less than $2.4 \cdot 10^{-9}$ by hypergeometric test. This indicates that the largest peak captures receptors that are important in angiogenesis. The statistical significance is higher than that for the distance from KDR and FLT1 calculated in the previous section, showing that the comberon is a useful concept for functional elucidation of receptor expression.

To quantify the significance of a receptor being included in the main peak (the main comberon) in the ordered dendrogram we performed a boostrap frequency analysis and a permutation test. See Table 1 for the genes with the largest bootstrap frequency. A hypergeometric test similar to that done above on the main comberon from Figure S4 was also done on the top genes from the mouse bootstrap shown in Table 1. In this set there are 7 known angiogenesis receptors, and this gives a p value of $5 \cdot 10^{-8}$. As negative controls we used 30,836 sets of genes from the GO Biological Process class (see Table S3). Table S3 shows that the only strongly significant associations are with terms related to angiogenesis.

The bootstrap frequencies of *Mus musculus* and *Homo Sapiens* data from PanglaoDB annotated by DCS and Alona and with an alternative DECNEO implementation using Alona are correlated, see Table S6, showing consistency among different methods. Table 1 and Table S4 show that the highest frequencies are obtained in the *Mus musculus*, as expected given the larger number of datasets.

Mus musculus		Homo sapiens	
Receptor	Bootstrap frequency	Receptor	Bootstrap frequency
ADGRF5	1.00	**CLEC14A**	1.00
CD93	1.00	**ESAM**	1.00
CDH5	1.00	**PECAM1**	1.00
ENG	1.00	**RAMP2**	1.00
ESAM	1.00	**ADGRF5**	0.98
FLT1	1.00	**ADGRL4**	0.98
KDR	1.00	**PTPRB**	0.98
PECAM1	1.00	**TIE1**	0.98
RAMP2	1.00	**FLT1**	0.95
ADGRL4	0.99	**CDH5**	0.91
PTPRB	0.97	**ENG**	0.91
TEK	0.94	**CALCRL**	0.89
TIE1	0.94	CAV1	0.84
CLEC14A	0.75	**CD93**	0.84
ROBO4	0.75	CD59	0.74
GPIHBP1	0.62	TGFBR2	0.73
FLT4	0.57	FCGRT	0.72
CALCRL	0.53	**KDR**	0.71
		THBD	0.71

Table 1. Main comberon bootstrap frequencies of PanglaoDB endothelial cells annotated by DCS, *Mus musculus* and *Homo Sapiens*, combination of measures. Only receptors where the values are above 0.5 in mice (this corresponds to p-value of less than 10^{-7}) and 0.7 in humans are shown. Receptors shown in bold with a yellow background are present in both lists and are therefore common to the main comberon of both species.

Receptors that have *Mus musculus* PanglaoDB-DCS "Combination of measures" bootstrap frequencies of at least 0.5 are included in Table 1. The complete frequencies data are available in Table S4. A permutation test (see Statistical analysis section of Methods) allowed us to obtain a non-parametric estimate of the probability of obtaining a frequency above 0.5 in *Mus Musculus* as p< 10^{-7} for column 1 of Table 1. A lower probability set (p<0.05) can also be identified with bootstrap frequency values between 0.33 and 0.5. This includes an additional 9 genes shown in column 1 of Table S4.

Statistical validation of the main comberon
Statistical validation was obtained by splitting the original Panglao dataset in two equal parts and by obtaining an independent dataset.
The results from the split datasets are shown in Table S2. The split was done in two ways, one where the splits had the same number of studies and the other where they had the same number of samples.

For the splits containing the same number of studies the Pearson correlation of bootstrap frequencies between the two splits was 0.87. The number of genes with bootstrap frequency above 0.5 in the first split dataset that was also found in the second split was 17 out of 23. This gives a p-value of $1.5 \cdot 10^{-27}$ using the hypergeometric test.

For the splits containing the same number of samples the Pearson correlation of bootstrap frequencies between the two splits was 0.91. The number of genes with bootstrap frequency above 0.5 in the first split dataset that was also found in the second split was 16 out of 26. This gives a p-value of $3.3 \cdot 10^{-24}$ using the hypergeometric test.

For the independent dataset we obtained 14 additional single cell mice gene expression studies, containing 71 samples and 65,057 endothelial cells (see Table S5). The bootstrap frequencies of the independent dataset were compared to those of the original dataset (first column of Table1). The comparison is shown in Table S12. The Pearson correlation of bootstrap frequencies between the original and the independent datasets was 0.85. The number of genes with bootstrap frequency above 0.5 in the main dataset that was also found in the independent dataset was 16 out of 25. This gives a p-value of $1.6 \cdot 10^{-27}$ using the hypergeometric test. The dendrogram for the independent dataset is shown in Figure S12.

Comberons from Multiple Peaks

In addition to the main peak analyzed in Table 1 multiple other peaks can also be detected. Using the "Combination of measures" of *Mus musculus* endothelial cells of PanglaoDB-DCS we identified the group with the highest bootstrap frequency and peak height. This group (group P1 in Table S7, blue cluster in Figure S9) is composed of the same 13 receptors which had the highest bootstrap frequency in Table 1 (mouse data). In addition to this main group of receptors we identified 10 other receptor groups (comberons), shown in Table S7. The names of the groups, P1 through P11, reflect ranking in their average bootstrap frequency.

We performed a Metascape analysis on the 11 groups and found that groups P1, P6, P9 and P11 are all enriched for blood vessel development, cell migration and growth, and peptidyl-tyrosine phosphorylation, see Figure S11. The detailed KEGG, GO and Reactome gene enrichment analysis of the 11 groups, Table S9, using PyIOmica (Domanskyi et al., 2019a) enrichment functions confirms the results of the Metascape analysis. The evidence for the biological functions of the minor comberons is however not as strong as for the main comberon.

Tissue-specific receptors

There are 8 *Mus musculus* and *Homo sapiens* tissues within the PanglaoDB database annotated by DCS for which there are substantial data (at least two studies with at least 300 endothelial cells each including at least 5 samples). The correlation among the 21 studies within these 8 tissues is shown in Figure 3. Most studies belonging to the same tissue cluster together. Additionally, the hypothalamus, olfactory bulb and subventricular zone studies, which are all parts of the brain, also cluster together. The average correlation values within tissues of the same type are almost always larger than the average correlation with other tissues (see Table S10). These results indicate that endothelial cells have tissue specific co-expressed receptors.

Figure 3. Analysis of endothelial cells from selected studies in PanglaoDB database annotated by DCS. Clustering of Pearson correlation of the largest peak bootstrap frequencies for the combination of measures. Studies from almost all similar tissues group together.

Figure 4 shows the receptors with the highest bootstrap frequencies (included in the comberons) for the 7 tissues shown in Figure 3 and for one addition human dataset obtained from the choroid of the eye (Voigt et al., 2019). These tissues usually also have high bootstrap frequencies for the top 13 receptors from the general analysis of Table 1 (receptors left of the red line in Figure 4).

Figure 4. Comberons of endothelial cells of specific tissues from PanglaoDB and choroid annotated by DCS. Under All we include the global analysis shown in Table 1 for comparison. The 13 receptors to the left of red line are the 13 top receptors from Table 1 (with boostrap frequency higher than 0.9). The rest of the receptors are within the top-10 tissue-specific genes for at least one set.

Two tissues, lung and choroid, were selected for further analysis using Metascape. To focus on tissue specific pathways, we did not include the generally relevant 13 genes shown on the left of Figure 4. In the lung we found the "cell-matrix adhesion" GO term as the top significantly enriched term (Table S11), with q-value of less than 10^{-6}. The choroid results are also analyzed more in depth in the next section. In the choroid, the Metascape enrichment analysis yielded blood vessel morphogenesis as a top significant GO term (Table S11), with q-value of less than 10^{-9}. By removing the receptors contained in the top enrichment term we found that the second most significant GO term for lung is cyclic-nucleotide-mediated signaling and for choroid is the JAK-STAT signaling pathway, (Figure S13).

Choroid endothelial receptors and their ligands

Given the heterogeneity between tissues and the presence of tissue-specific receptors in the comberons we analyzed tissue-specific data to interpret and possibly guide experimental investigations. We futher analyzed the main comberon of human choroid endothelial cells, also in the light of very recent experimental results we have obtained (Li et al., 2021). The choroid vessels of the eye play an important role on the development of Age-related Macular Degeneration (AMD) one of the main causes of blindness (Li et al., 2021), Several VEGF inhibitors are used as FDA-approved therapies for AMD but not all forms of the diseases benefit from them (Li et al., 2021).

In Table 2 we show the 25 receptors with a bootstrap frequency higher than 0.7 in human endothelial cells (see also figure S14). Single cell RNAseq data were from Voigt et al (Voigt et al., 2019). Out of these 25 receptors, 6 have the strongest evidence of relevance to choroid vasculature, either because they are targets of approved AMD drugs (KDR and its co-receptor NRP1), are associated with increased AMD risk in GWAS studies (CFH) (Fritsche et al., 2016) , or have been shown to regulate choroid

angiogenesis in our separate experimental paper (Li et al., 2021) (LIFR, OSMR, IL6ST). Out of the remaining 19 receptors, 16 have at least 10 publications associating them with angiogenesis in a systematic Pubmed search.

Receptors	Frequency	Pubmed angiogenesis	Strongest Experimental evidence
BMPR2	0.9	92	
PTPRG	0.9	3	
PTPRB	0.89	9	
NOTCH1	0.87	388	
PTPRM	0.86	0	
ENG	0.85	891	
IL6ST	0.82	13	LIFR and OSMR co-receptor
LIFR	0.82	8	Li et al LIF paper
NRP1	0.8	399	KDR co-receptor
PLXND1	0.78	31	
ITGA5	0.77	18	
OSMR	0.75	9	Li et al LIF paper
PECAM1	0.75	404	
CD46	0.74	22	
CDH5	0.74	51	
KDR	0.74	1952	main VEGFA receptor
ROBO4	0.74	74	
TEK	0.74	210	
TGFBR2	0.74	194	
TIE1	0.74	205	
ADIPOR2	0.73	11	
S1PR1	0.71	68	
CFH	0.7	69	Main genetic risk factor for AMD
IL4R	0.7	13	
ITGAV	0.7	26	

Table 2. The receptors in the main comberon of human choroid endothelial cells. All the receptors with substantial evidence for a role in angiogenesis are shown in red.

Furthermore, 12 of the receptors of the choroid main comberon have been previously shown to intereact, either as co-receptor pairs, or as sets with common functional roles (Frye et al., 2015; Tzima et al., 2005) or as receptors for ligands showing synergy in stimulating choroid endothelial cells (LIF and VEGFA) (Li et al., 2021) (Table 3).

Co-receptor pairs				
TGFBR2	ENG			
KDR	NRP1			
LIFR	IL6ST			
PLXND1	NRP1			
OSMR	IL6ST			
Other known interacting receptors				
PTPRB	TEK	CDH5	KDR	Frye et al
KDR	PECAM1	CDH5		Tzima et al
LIF (LIFR)	VEGFA (KDR)			Li et al

Table 3. Multiple receptors in the main comberon of human choroid endothelial cells are part of known interacting pairs or sets.

Additional evidence for a natural role of the choroid main comberon receptors is provided by the finding that several of their ligands are co-expressed in human choroid fibroblasts and significantly associated with VEGFA in the co-expression dendrogram (see Figure 5)(p<0.05). The signal is not associated with the main ligand comberon in choroid fibroblasts but with a minor one (Figure S15) which is consistent with the fact that angiogenesis signaling is only one of several possible roles for fibroblasts. A significant association with VEGFA (p<0.05) is however also found in the PanglaoDB-Alona mouse fibroblasts, when looking at ligands for the known angiogenesis receptors (Figure S16).

These results provide initial insight into the communication mechanisms among different choroid cell types.

Figure 5. Dendrogram of ligands co-expression in human choroid fibroblasts. Ligands that are co-expressed in individual cells are closer to each other in the dendrogram order. Ligands for the choroid endothelial comberon receptors are labelled under the dendrogram. The red pointer indicates VEGFA, the main KDR ligand.

DISCUSSION

Our results show that single cell data can help us to identify a set of co-expressed endothelial cell receptors, conserved among species (mice and human) and enriched, within a network, of connections to up-regulated genes. This set includes VEGF receptors and also includes several other receptors previously shown to play a role in angiogenesis. We can combine endothelial single-cell data from multiple organs and show organ-specific receptor profiles in these naturally occurring combinations.

Many advances have contributed to increase the understanding of the molecular machinery coordinating gene expression (Komili and Silver, 2008; Latchman, 2015), but we know less about the functional effects of combinations of expressed genes in the human body. Our method identifies combinations of receptors by integrating co-expression in single cell with 3 additional measures: expression fraction, network enrichment and evolutionary conservation. Importantly, these 3 measures do not need to be increased in all members of a comberon, so that the method is able to deal with incomplete information and with evolutionary change. The procedure has led to sets including clinically validated drug targets, like the VEGF receptors, and to receptors identified in the choroid in recent experiments (Li et al., 2021). This integrated computational and experimental approach provides an effective strategy for the study of comberons.

The conclusions are supported by multiple highly significant statistical tests obtained from large datasets, including many independent experiments. Statistical validation was reinforced by splitting the original dataset in two parts, each analyzed separately and by obtaining an additional independent set of single cell gene expression data. Evolutionary pressure to preserve functionally important and biologically robust combined effects provides a plausible mechanism for the evolutionary conservation of the main endothelial receptor comberon between mice and humans. The finding that endothelial receptors work in combination is consistent with the observation that drug combinations are needed to improve clinical response in angiogenesis (Apte et al., 2019; Chung and Ferrara, 2011; Ferrara and Adamis, 2016).

We have also adopted a strategy of extensive internal replication to further validate our findings. Two versions of the software implementation of the algorithms have been written independently by two of the authors (SD and AH), to guard against coding errors and unforeseen consequences of implementation choices. We have also used two different methods to identify endothelial cells in single cell RNAseq data: DCS, (Domanskyi et al., 2021; Domanskyi et al., 2019b) which was previously developed by some of us, and Alona (Franzén and Björkegren, 2020). Other validation strategies also represent replications in a broader sense. To identify evolutionary conservation, we studied similarities in comberons among two species, mouse and humans. We also found in specific tissues many of the receptors that were present in the global comberons. The strategy used to identify comberons (called Combination of measures in Figure 2) showed a strongly significant enrichment for the presence of receptors already known to play a role in angiogenesis.

Recent papers have started to talk about the syntax of intercellular communication to indicate how combinations of signals are used by cells (Rieckmann et al., 2017). The word syntax, from the ancient Greek "arrange together", is also used in linguistics and in

computer coding with an analogous meaning for signals of a different type. The study of syntax in these fields has a longer history and advanced quantitative models have been developed (Boeckx and Grohmann, 2013; Köhler, 2012). Using the same word might help scientists from these fields to contribute their expertise. The study of intercellular communication syntax by Rieckmann et al (Rieckmann et al., 2017) used bulk measurements and they were able to investigate combinations of signals used by different cell types. The approach used here is based on single-cell measurements and it allows progress towards the identification of signals used by a given cell type in a particular cell state, in our case a pro-angiogenic state.

The Discussion has been written with the contribution of a co-author with a background in the Economics of innovation, who also contributed to the statistical analysis. This has led to the use of language more accessible to non-specialist audiences and it is also an attempt to address the need for the "public engagement with science" which has been stressed by professional societies (Agre and Leshner, 2010) and funding agencies (e.g., the National Science Foundation's Broader Impacts Criterion).

Biological implications

In many scientific disciplines it is normal to examine the results of an experiment in the light of general models. For our biological results we can to refer to a recent description by Paul Nurse of the five major ideas that underpin Biology (Nurse, 2020). Paul Nurse received the Nobel Prize for his cell biology discoveries, and he is the Director of the Francis Crick Institute. He has stated that the ideas he has presented are not new but that even for an experienced scientist like him putting them together was illuminating (Nurse, 2020). Our work has interdisciplinary implications and it is especially helpful for non-specialists to see the results from a wider perspective.

The five fundamental ideas of Biology according to Paul Nurse are the cell as the basic unit of life, the gene as the basis of heredity, evolution by natural selection, chemical activities bringing about life and life as managing of information. The approach described here touches on all of them, including single cell measurements of gene expression, evolutionary conservation, molecular interactions of receptors and biological network models of information transmission and integration. It is therefore not surprising that the approach would eventually prove informative, but it is still important to show that available datasets can already provide a clear signal.

If a set of co-expressed receptors is evolutionary conserved it is reasonable to ask why. The main set we have identified (see Table 1) includes VEGF receptors and several other receptors clearly involved in angiogenesis. In mice, where more data are available, it includes at least 18 receptors. This redundancy is common in biology and has been shown to lead to robust control (Feala et al., 2012; Nowak et al., 1997; Wagner, 2005). It might also be byproduct of the tinkering action of evolution, which often proceeds by gene duplications and can lead to not fully adaptive solutions (Jacob, 1977; Ohno, 2013). A classic example is the blind spot of the human eye, which is consequence of the inverted retina, not present in the octopus (Dawkins, 1997).

Many of the receptors in this set are known to promote angiogenesis however some (like PLXND1 in the lower probability set (Sakurai et al., 2010; Yang et al., 2015) and choroid and OSMR in the choroid) are considered to be angiogenesis inhibitors. This is not surprising if we consider that biological regulation aims to obtain homeostasis and

controlled response, rather than uncontrolled activity (Nurse, 2020). In fact, it has been already reported that MEF2C, a transcription factor induced by VEGF, inhibits angiogenesis, providing negative feedback (Sturtzel et al., 2014). It is also possible that the expression of positive and negative regulators on endothelial cells might assist in guiding the direction of vessel growth. This is an example of possible insights into the syntax of intercellular communication that can be provided by single cell but not by bulk data. Using bulk data we would not be able to establish if molecules with opposing functional roles were expressed by cells in different conditions or within the same cell.

Pharmacological and medical implications

As we have mentioned in the Introduction, the problem of finding anti-angiogenic drug combinations is still largely unsolved, as it is also for many other complex diseases and conditions. Traditionally drug targets are individual proteins but mimicking the physiological control of cells would suggest that a comberon could be the appropriate drug target for complex diseases. It is not unfeasible to target receptor comberons, because Ab drug inhibitors for each receptor can be developed. There are around 1,000 human cell membrane receptors (Almen et al., 2009) (excluding olfactory receptors and other specialized types) and until now Ab drugs have been approved by the FDA against 26 of them (Antibody Society, 2020). Our work supports the substantial acceleration of Ab drug development against receptors and their ligands, because it suggests a new way to use these drugs, initially as research tool and later as therapies. While in vitro models, for example organoid models using iPS cells, can be informative, they can only be stepping-stones to clinical studies if we want to elucidate human intercellular communication *in vivo*. The last step can only be achieved with therapeutic-grade antibodies.

The results shown here are clearly only one more reason for Ab drug development acceleration, in addition to several others, including the increasing role of Ab as individual drugs (Strohl, 2018), the prospect of advances in decoding of intracellular signaling, possibly using methods similar to those used to decipher encrypted messages (Singh, 1999) and developments in the use of artificial intelligence to predict effective drug combinations (Kuenzi et al., 2020).

The process used by Shinya Yamanaka and his group to identify combinations of factors to induce pluripotent stem cells can be used as an informative analogy for therapeutic interventions on comberons (Karagiannis et al., 2019). Starting from a longer list of transcription factors he used several methods, including gene expression, to identify a set of 24 transcription factors that he then reduced experimentally, finding that just 4 of them were sufficient (Karagiannis et al., 2019). Even if receptor comberons are quite large, biological systems respond dynamically and acting on just some of the receptors might be sufficient to induce a stable change from pro to anti-angiogenic activity. It is likely that this last step will need to be conducted experimentally as was the case for the Yamanaka factors.

Beyond angiogenesis, it is clear that important cellular processes are mediated by a large number of stimulatory and inhibitory signals mediated by membrane receptors. Different cell types form complex and in part organ-specific intercellular signaling networks (Armingol et al., 2020; Rieckmann et al., 2017). For example, the response to injury and infection is mediated not only by the cells of innate and adaptive immunity but

also by other interacting cell types like endothelial cells, epithelial cells and fibroblasts (Gomes and Teichmann, 2020; Krausgruber et al., 2020). Our understanding of these interactions is however incomplete, and our therapeutic tools are limited and do not match the complexity of natural signals.

It has been noted that the past decade has witnessed unprecedented pandemic explosions and that we have entered a new pandemic era, due to a complex set of causes including human population growth, increased mobility and environmental factors (Morens and Fauci, 2020). The COVID-19 pandemic has stimulated our understanding of the importance of controlling the response of endothelial and immune cells to infectious agents. For example, Ackerman et al (Ackermann et al., 2020) have shown that COVID-19 patients exhibit severe endothelial injury and perivascular T-cell infiltration. An imbalanced host immune and inflammatory response drives the development of COVID-19 (Blanco-Melo et al., 2020). The randomized RECOVERY trial (The RECOVERY Collaborative Group, 2021), showing the beneficial effect of dexamethasone on COVID-19 mortality, provides perhaps the strongest evidence about the clinical advantage of strategies acting on the immune response, which we cannot however yet calibrate precisely. Developing pharmacological tools to control the signals among cells of the immune and inflammatory system is an aspect of pandemic preparedness that might be deployed as part of a fast response, before virus-specific drugs and vaccines are available.

Social and economic implications

While developing Ab drug is feasible, a clear barrier for a biomimetic strategy targeting receptor comberons is the very high and increasing cost of drug development (Scannell et al., 2012). Combining drug development with fundamental scientific understanding of *in vivo* human intercellular communication is challenging. As a research area that combines elements of basic platform science and its application to therapeutics for specific condition, it does not fit neatly into the distinct priorities of the two largest sources of US biomedical funding, federal support (mainly NIH) and industry.

A recent report from the National Academy of Science (National Academies of Sciences, 2018) pointed out that the benefits of some innovative projects have previously been accepted slowly due to resource limitations. They state that "much of the biomedical research community was strongly opposed to the Human Genome Project when it was first proposed, believing that it diverted resources from more valuable investigator-driven work (Palca, 1992). The project and its impact look much different in hindsight." (National Academies of Sciences, 2018). Diverting public resources from valuable scientific programs is however a legitimate concern. The academic community is also not equipped to take drugs all the way to approved clinical products and is not likely to be supported by NIH to do so.

Industry funding might also be only partially appropriate. The fundamental nature and the complexity of the problem of intercellular communication within organisms, including human beings, would benefit from full scientific openness and sharing (National Academies of Sciences, 2018). This complete openness is not possible in the pharmaceutical industry especially at the preclinical stage. Even recent high-profile papers from pharmaceutical scientists showing that many academic studies are not reproducible were not able to specify the studies they were referring to (Begley and Ellis,

2012; Prinz et al., 2011). This is an obstacle to the self-correcting mechanisms of science. Evidence that Celera's gene-level intellectual property hindered subsequent research investments and innovation offers a cautionary tale in this regard (Williams, 2013). The case of genetically-modified mice also suggests that openness can increase innovation through greater exploration of novel research directions (Murray et al., 2016). It is precisely this type of exploration that is more likely to yield breakthrough innovations (Azoulay et al., 2011). The key to ensuring a platform that can enable a wide range of scientific discoveries and therapeutic advances is to place it in the public domain such that would be innovators can access that knowledge with few barriers. The platform will be composed of approved Ab drugs for most receptors and, giving the increasing cost of individual drugs (Scannell et al., 2012), using only industry resources is likely to make the cost of large combinations unsustainable.

The best opportunity is provided by large philanthropic sources of funding, which are able to combine non-profit motives with industrial collaborations. We refer specifically to the increasing number of billionaires joining the Giving Pledge (Giving Pledge, 2020). These philanthropists (211 at present) have committed to donate the majority of their wealth. According to their statements (Giving Pledge, 2020) they are looking for unique projects with a potential large impact on society that are not supported by other sources and many also wish to contribute their entrepreneurial skills. The Giving Pledge was started in 2010, however a recent report of the Institute for Policy Studies (Institute for Policy Studies, 2020) has found that the total wealth of most of the Pledgers has increased since 2010. Some have however provided large donations and there is also an example, that of Charles Feeney, who has donated his entire wealth (O'Clery, 2013). Even using a conservative estimate the total of available funds is larger than ten times the annual budget of the NIH, the biggest supporter of academic biomedical research in the US (Institute for Policy Studies, 2020). The shortage of suitable projects with a potential large impact that Pledgers would support could be addressed by one or more joint initiatives of the biomedical community.

Cell communication is a fundamental unsolved biomedical problem and Ab drugs could be provided as a platform technology to academic and industrial groups, resetting the traditional border between basic science and drug development. Providing tools rather than directing research is consistent with many previous examples of scientific innovation suggesting that a diversity of approaches should be encouraged (Azoulay et al., 2011). The efforts of other scientific communities that have been able to define priorities for large-scale projects might be informative. An example is the Snowmass process (Cho, 2020) which has been used by particle physicists for several decades and has guided the choice of major projects after encouraging discussions and feedback from all scientists with relevant expertise. The Snowmass process aims to produce a scientific consensus, as encouraged by federal funding agencies, especially the DOE (Cho, 2020). An alternative that was able to facilitate and record the expression of a large diversity of opinions in addition to the preferred choices, using dedicated technology, is provided by a recent survey of more than 4,000 biomedical scientist (Maxmen, 2020). The scientific community could greatly promote human health and the efficient use of large philanthropic resources by providing an open critical evaluation of broadly beneficial endeavors similar to that outlined here.

General scientific implications and limitations

While publication is clearly a very useful step, this project also needs to adopt a format suitable for continuous updating, to keep up with the increasing amount of available single cell data. These data are likely to keep growing at even faster rate, considering that comberon identification can be extended to other cell types, go beyond receptors and include more species to measure conservation.

We have only investigated the more general properties of receptor comberons in endothelial cells. For example, some receptors might belong to more than one comberon. A first step in this direction is provided by the results we present on tissue-specific receptor analysis, but there might be comberons associated with different cell states within a tissue, including disease states. More data will certainly improve the precise identification of the receptors included in the comberons, especially at the organ and tissue-specific level. It will also be interesting to use the tools we describe to study how aging affects intercellular communications. Recent data show, for example, that angiogenic signals from fibroblasts to endothelial cells change with aging (Kaur et al., 2016).

It is important to promote scientific replication, not only internally as we have done here, but also externally. We facilitate this by sharing our software code in an organized format.

Conclusions

This paper demonstrates the existence of evolutionarily conserved and naturally occurring combinations of receptors in endothelial cells. We also show tissue-specific combinations and, in the case of choroid endothelial cells, consistency with well-established and recent experimental findings.

MATERIALS and METHODS

Data sources

Three sources of data were used in our analysis:

(1) The primary data source in our study was PanglaoDB, a web-database of scRNA-seq datasets of *Homo sapiens* and *Mus musculus* studies, that collects 22,382,464 cells from 236 studies (SRAs, or Sequence Read Archive submission accessions) in 1368 samples (SRSs, or Sequence Read Archive sample accessions), spanning numerous tissues and cell types (see Table S5). After quality control PanglaoDB provides data from 4,459,768 *Mus musculus* cells and 1,126,580 *Homo sapiens* cells. Each of the datasets were downloaded from NCBI SRA and mapped to the GRCm38/GRCh38 genome assembly by the authors of PanglaoDB (Franzén et al., 2019). All cells in PanglaoDB are annotated with cell type labels by Alona (Franzén and Björkegren, 2020), a web-based bioinformatics tool for scRNA-seq analysis. This cell type annotation is incorporated in the PanglaoDB database. In addition to the Alona annotation, we processed each of the datasets in PanglaoDB using a platform for cell type annotation that we have recently developed (Domanskyi et al., 2021; Domanskyi et al., 2019b), DCS 1.3.6.10.

(2) We also obtained an independent dataset for validation, composed of 14 single cell mice gene expression studies, containing 71 samples and 65,057 endothelial cells (see Table S5). This independent dataset was obtained after the analysis of the main dataset had already been completed and did not affect any of the methodological choices.

(3) Another source of data was a dataset from Voigt et al. (Voigt et al., 2019) containing *Homo sapiens* choroid scRNA-seq on 18621 cells in 8 samples enriched for endothelial cells. The dataset was downloaded from NCBI and re-mapped with CellRanger 4.0.0 and the GRCh38 genome assembly; each of the 8 samples has 4 paired-end reads sequencing runs (SRRs) that were merged into one. In our alternative implementation of DECNEO, used as a validation, we used the original Voigt et al. (Voigt et al., 2019) data mapping, quality control, and cell type annotation of the choroid cells.

Input data format

See Figure 6 for a graphical summary of the analysis pipeline. The method is designed to take two single-cell transcriptomics datasets corresponding to two different species, e.g. *Mus musculus* and *Homo sapiens*. Each of the input datasets has to be in a table format where gene identifiers are in rows and cells and samples identifiers are in columns. Where possible, gene identifiers should be converted to one species, e.g. *Homo sapiens*, using their homologs. Starting from the gene counts matrix, the gene expression values should be scaled so that the sum of values of each cell in the dataset is identical (we recommend a scaling factor of 10^5) and then log-transformed. Cell type and sample annotation is required for each cell in each dataset, assigning cells to two different classes, e.g. endothelial cells and non-endothelial cells. The preprocessing step described above is necessary for the analysis pipeline. The preprocessing step is not directly included in DECNEO, but DCS (Domanskyi et al., 2021; Domanskyi et al., 2019b), can be used to normalize and annotate cells.

Figure 6. Summary of the analysis pipeline. Normalized single cell transcriptomics datasets are processed in steps depicted in panel A. The principal steps of the pipeline are outlined in panels B, C and D. An analysis report is generated at the end of the pipeline.

Dataset validation and bootstrap resamples preparation

First, the dataset is validated and checked for sufficient number of samples (see Figure 6). If the number of samples in a dataset is below 5 we combine them and then split the data into 10 non-overlapping sets, to allow bootstrapping. Next, within each sample $= 1 \dots k$, we calculate a matrix $\widetilde{M}_i$ of distances between gene pairs based on gene expression according a user-specified metric (Euclidean, cosine, or correlation distance). In a given sample, only genes that are expressed in at least 5% of cells are considered. For the choroid dataset the cutoff is 1% because of the high quality of this dataset, which was experimentally enriched for endothelial cells and had a higher sequencing depth (22). For each sample, the distance is calculated between each gene in a set of m genes of interest (in Figure 2 the genes of interests are all known receptors) and each gene in the complete set of n expressed genes, resulting in M, an $n \times m$ matrix that contains the n genes in rows and m receptors in columns. Samples (ie, the set of k matrices, one per sample) are then aggregated using the median distance among non-missing values. For missing values in the distance matrix M we assume no correlation

or, in the Euclidean case, replace with the maximum of M. When the matrix M is calculated using all the available samples, we refer to this as the original data. In addition to the *original data*, we prepare a set of *bootstrap resamples* by resampling with replacement all samples N times (we used $N = 100$): i.e. we randomly choose with replacement k samples and define a distance matrix following the same procedure as in the *original data*. The size of the bootstrap resamples was validated for one dataset, the choroid dataset, where we compared the results of 100 and 1000 resamples. The output of the bootstrap was very similar among the two independent resamples. The bootstrap frequency Pearson correlation between them was 0.99. The set of bootstrapped distance matrices is used to analyze the effect of gene expression variability and for statistical inference, as explained later in the statistical analysis section.

Dendrogram ordering by co-expression

The ordering of the dendrogram (see Figure 2) is based on the Pearson correlation of normalized single cell gene expression, which is a commonly used gene co-expression measure. The similarity metric is the Euclidian distance of the Pearson correlation. We also compared alternative choices for these metrics, for example we tested Spearman and Pearson correlations of gene expression (see Table S1).

Here, an $m \times m$ symmetric matrix A with no missing values is calculated from each $n \times m$ matrix M by considering the Euclidean distances of the Pearson correlations between the columns of M, thus defining a distance matrix for the m receptors only. We then perform an agglomerative clustering on the matrix A using a Ward variance minimization linkage method (Ward Jr, 1963) and construct a dendrogram (alternative clustering similarity metrics and linkage methods are compared in Table S1). Correlation is compared to the Euclidean distance of Pearson correlation in the dendrograms shown in Figures S1, S2 and S3. Finally, receptors, which are the leaves of the dendrogram, are ordered to minimize the sum of the distances between adjacent leaves in the dendrogram layout (25). This dimensionality reduction procedure defines a "dendrogram distance" between any two ordered receptors, which is a positive integer equal to the difference of their ordering index. The N bootstrapping resamples and the original data are processed according to the steps shown in Figure 6.

Following this procedure genes that are co-expressed in individual cells are closer to each other in the dendrogram order.

The 3 measures shown in the lower half of the Dendrogram Figure (Figure 2) are:

1- Expression fraction

For each receptor we calculate the fraction of cells that express that receptor at non-zero level. We calculate the median across the samples. We expect important receptors to be expressed in most cells.

2- Network enrichment

To find receptors close to genes that are differentially expressed in endothelial cells we use a predetermined undirected gene interaction network known as *Parsimonious Composite Network* (PCN) (Huang et al., 2018). This network contains 19,783 genes and was obtained by combining data from 21 popular human gene

interactions networks. Only interactions supported by a minimum of two networks were included. The 21 networks included interactions derived from multiple methods, such as protein-protein interactions, transcriptional regulatory interactions, genetic interactions, co-expression correlations and signaling relations. The authors performed an extensive evaluation of the PCN network showing that it was superior to any of the 21 networks used to build it. We independently determined that sets of genes with similar functions are close to each other within this network. Using inflammation, cell cycle and metabolic gene sets we measured intraset efficiency, a metric we have previously introduced (Ong et al., 2015), and the enrichment for direct connections between the genes in a given set using the binomial distribution. All the tests gave highly significant p-values, supporting the high quality of the PCN network.

First, we determine the set of the top 1000 differentially expressed genes, S, in a dataset consisting of several samples. The differential gene expression for each sample is computed using a t-test between gene expression of endothelial cells and non-endothelial cell in that sample. Genes with a p-value below 0.001 are retained and ordered according to the decreasing t-test statistic. The median of the ranks across samples is taken to rank the differentially expressed genes of the entire dataset.

The t-test was also compared with the Mann-Whitney U test and the network enrichment obtained was very similar (see Table S1 for bootstrap frequency obtained with the two methods).

To obtain a p-value we use the binomial distribution. Let p be the probability that a randomly chosen PCN edge is connected to a gene in the set S of differentially expressed genes. We estimate p as the number of edges with at least one gene in S divided by the total number of edges in the network. For each receptor j the probability to have k_j or more enriched edges out of its total degree n_j in PCN is given by

$$P_j(k_j) = 1 - \sum_{i=0}^{k_j-1} \binom{n_j}{i} p^i (1-p)^{n_j-i}.$$

The negative logarithm in base ten of this probability is the "Network enrichment" measure used in our method.

This measure reflects the level of expression of the first neighbors of a receptor within the PCN network. More exactly, it captures the enrichment in number of connections of a receptor to the set of differentially expressed genes.

3- Evolutionary conservation

To evaluate evolutionary conservation, we calculate the receptor dendrogram ordering using data from two different species and we quantify their similarity. The comparison is done by pairing data from different species. In this paper we use *Mus musculus* and *Homo sapiens*. For each receptor j in the first case (for example *Mus musculus*) we count the number of receptors in a certain window in the dendrogram ordering (here of size 50, centered at j) that are also present in a window of the same size centered at homologous receptor j for the second case (for example *Homo sapiens*). The number of overlapping receptors in the neighborhoods of the receptor j defines a local measure of evolutionary conservation. Genes that do not appear in both species are marked as missing with a grey bar.

Other measures

We have considered other measures, including evolutionary rate and age, average gene expression, distance between neighboring leaves along the dendrogram, evolutionary conservation of known angiogenesis receptors, and known angiogenesis receptors from gene ontology angiogenesis terms. These measures did not help in identifying the comberon or in the enrichment for known angiogenesis receptors. The main implementation of DECNEO allows for the inclusion of these measures, or even new, user-defined, measures in future analysis.

Known angiogenesis receptors

We use for validation a set of 22 receptors that are known to regulate angiogenesis: CD36, CD47, CLEC14A, CXCR2, ENG, FGFR1, FGFR2, FGFR3, FLT1, FLT4, KDR, KIT, MET, NRP1, NRP2, PDGFRA, PDGFRB, PLXND1, ROBO1, ROBO4, TEK and VLDLR. These receptors were obtained before the analysis from a systematic Pubmed search for receptors with functional effects on angiogenesis. The list includes FLT1 and KDR, the main VEGF receptors, and also other targets for anti-angiogenic drugs under development (Chung and Ferrara, 2011; Ferrara and Adamis, 2016; Kong et al., 2017; Potente et al., 2011). As negative controls we used 30,836 sets of genes from the GO Biological Process class (see Table S3).

Combination of measures and comberons

Figure 2 shows how we analyze the combinations of different measures. The lowest panel in Figure 2 show a linear combination of the measures. "Combination of measures" is a sum of "Expression fraction", "Evolutionary conservation", and "Network enrichment".

For each dataset d and for each resample c, we first normalize these measures so that the sum of each measure is one and take a moving average along the dendrogram to smooth the normalized measures (using a centered window of 21 receptors and a uniform averaging kernel). Then, we add their values and take a moving average of the combined measure to obtain a combination score m_c^d.

In the analysis, we identify the highest peak p_0 in m_c^d, such that p_0 is $\arg\max(m_c^d)$, along the dendrogram and we find all the receptors with $m_c^d \geq 0.5 m_c^d(p_0)$ that are in the immediate proximity of this main peak. Then, we count how many times a given receptor is found in the main peak across all bootstrap resamples in a dataset and we calculate its frequency of peak inclusion. The software can also calculate the frequency of inclusion for the main and all other secondary peaks, p_i where $i \geq 0$, in the dendrogram by considering all extrema of height and prominence larger than 0.05 separated by at least 50 receptors. Prominence of a peak is defined as the vertical distance between the peak and the lowest contour line encircling it but containing no other higher peaks. When secondary peaks are used, each peak i (main or secondary) contains at most 50 receptors that are in the immediate proximity of the peak and have $m_c^d \geq 0.5 m_c^d(p_i)$, see Figures 2 and S4.

The sets of genes that are included in the peaks are the comberons. It is important to observe that the measures used to identify the peaks are not necessarily increased for

a particular gene. For example, as can be seen from Figure 2, a gene that has a low value of network enrichment can be included in a peak because is co-expressed with other genes that have a high value for this measure.

Visualization

A "plot" mode (see Figure 6C and in part Figure 2) can be selected to visualize dendrograms, heatmaps of the A matrix (as an option), and several other panels. These panels plot the measures described above, such as the Evolutionary Conservation and Expression Fraction, smoothed using a moving average along the dendrogram (indicated by an orange line), the combinations of these measures and highlight the location in the dendrogram of the known angiogenesis receptors.

Statistical analysis of Dendrogram Ordering

The likelihood-ratio test was used to analyze the distance in the dendrogram of known angiogenesis receptors from KDR and FLT1 . The null hypothesis here is that the known angiogenesis receptors are uniformly spread in the optimally ordered dendrogram, whereas the alternative hypothesis is that the known angiogenesis receptors should be in the vicinity of KDR (or FLT1) in the dendrogram, showing that known angiogenesis receptors are co-expressed. Under the null hypothesis every dendrogram receptor has equal probability to be one of the known angiogenesis receptors. With the alternative hypothesis we use a logistic regression to get the probabilities of the receptors to be a known angiogenesis receptor given the optimal dendrogram ordering. The p-value is calculated using the survival function of chi-squared continuous random variable.

An alternative non-parametric method was also used for the same purpose (Wooldridge, 1999). The robust Poisson pseudo maximum likelihood estimation (PPMLE) was implemented in R using the package *gravity* (Silva and Tenreyro, 2006) . The p-value was obtained from the difference of the deviances of the null hypothesis and alternative hypothesis using the chi square distribution.

Statistical analysis of Comberons

The main aim of this statistical analysis was to show that there is a set of receptors significantly associated with the largest peak of Figure 2. We call the set of receptors associated with the largest peak the main comberon.

The first step was the use of the bootstrap method (Efron and Hastie, 2016) to improve the robustness of the dendrogram. This is a very common procedure in many scientific fields where dendrograms are used (Soltis and Soltis, 2003). In fact, the paper describing the application of the bootstrap method to dendrogram analysis is one of the 100 most highly cited papers of all time (Felsenstein, 1985; Van Noorden et al., 2014).

The bootstrapping involved 100 resamples with replacement. The values shown in Table 1 are bootstrap frequencies where 1 means that the receptor was found in the largest comberon in all of the 100 resamples. See Table S4 for the complete list of the bootstrap frequencies.

We then performed several statistical tests:
1- A permutation test (also called randomization test) (Good, 2013; Manly, 2006) for the co-expression that was used to build the dendrogram. We wanted to determine the statistical significance of the frequencies in the first column of Table 1, which contains the

Mus musculus data analyzed with DCS, and we therefore used a control where the ordering of the dendrogram was no longer dependent on correlation of gene expression of the receptors. First, we determined the maximum frequency of genes included in the main peak for 100 instances (to match the bootstrapping procedure) with randomized ordering of the dendrogram shown in Figure 2A. We then repeated this procedure for 10^7 realizations. This calculation was done using the High Performance Computing Center (HPCC) supercomputer at Michigan State University. From the 10^7 realizations we obtained a distribution of the maxima of the bootstrap frequencies for the original data case. We did not observe a frequency above 0.5 in any of the 10^7 realizations. The maximum observed frequency (0.49) was slightly below this number as shown in Figure S5. This procedure allowed us to obtain a non-parametric estimate of the probability of obtaining a frequency above 0.5. Note that this is a conservative estimate of the p-value given, because it is an estimate for a single value being larger than 0.5 while in first column of Table 1 there are 18 genes above this threshold. We also empirically determined that this distribution resembles a generalized extreme value distribution, which is a known distribution of maxima (see Figure S5).

2- As a validation we split the dataset in two equal parts. This was done twice, once splitting the samples and once splitting the studies. The bootstrap frequencies of the two splits were then compared using the hypergeometric test.

3- The significance of enrichment of the comberon for known angiogenesis receptors was calculated using the hypergeometric test. This was done for both the original and the bootstrap frequency of at least 0.5.

4- We also obtained an independent dataset composed of 14 single cell mice gene expression studies, containing 71 samples and 65,057 endothelial cells (see Table S5). The endothelial cells were identified using DCS as was done for the original Panglao dataset. This independent dataset was obtained after the analysis of the main dataset had already been completed and did not affect any of the methodological choices. The evolutionary conservation was obtained comparing the independent mice data with the human data from PanglaoDB. The bootstrap frequency of the independent dataset was compared with that of the original dataset using the hypergeometric test.

Implementation of the algorithm

The pipeline shown in Figure 6 is implemented in our new software DECNEO (DEndrogram of Co-expression with Evolutionary Conservation and Network Enrichment Ordering). We have developed a main and an alternative implementation of DECNEO that are open-source and maintained at https://github.com/sdomanskyi/decneo. Releases are published on the Zenodo archive.

The main implementation can be installed from PyPI using the command "pip install decneo". DECNEO modules are relying on Python numpy, scipy, pandas, scikit-learn, matplotlib and other packages (see Figure S6 for dependencies). The main implementation can leverage on a multicore large-memory High-Performance Computing Systems for fast processing of large datasets. The detailed documentation and installation troubleshooting instructions are available at https://decneo.rtfd.io/. The

processed data and results reported in this paper are deposited at https://figshare.com/. The source code of the alternative implementation is deposited at https://github.com/sdomanskyi/decneo/validation. The software was optimized for use with the HPCC supercomputer at Michigan State University but also tested on different operating systems and hardware.

Alternative implementation

The main implementation of DECNEO was validated by comparing its results to an additional implementation created independently as a software and methodology validation. The alternative implementation has allowed us to verify the robustness of the results under minor differences in the software implementation. Specifically, the differences are: (1) When finding differentially expressed genes for the network enrichment, the alternative implementation filters out genes that are downregulated and sorts them by p-value while the main implementation sorts genes by the median rank of the t-test statistic. (2) When computing the expression fraction for a set of samples, the median of the expression fraction across samples is taken, instead of the mean across samples as in the main implementation. (3) Lastly, in the alternative implementation, gene measures are not smoothed before merging them into a combination of measures, and the combination of measures is normalized so that the minimum is zero and the maximum is one. (4) In the alternative implementation only endothelial cells, but not non-endothelial cells, are used when generating pseudo samples.

The analysis of endothelial cells in PanglaoDB-Alona dataset is summarized in Figure S7 and the alternative implementation is shown in Figure S8. These analyses show results similar to Figure 2. Table S6 shows the correlation of bootstrap frequencies between DECNEO main and DECNEO alternative implementation.

Revealing Multiple Peaks

We observe multiple peaks in the measures and their combinations. These represent the main comberon and other minor comberons. See for example Figure S4, which corresponds to "Combination of measures" of *Mus musculus* endothelial cells of PanglaoDB-DCS. We analyzed all the peaks, including the largest one. Each instance of the 100 bootstrap resamples has a set of peaks with the receptors identified in the same way as in Figure S4. We order receptors (vertical axis) and peaks (horizontal axis) according to the "Ward" agglomeration method with Euclidean metric of binary membership, see Figure S9, which is colored according to each peak relative height. The left and top dendrograms highlight receptors and peaks groups. To focus on the most evident peaks we obtained 15 groups of receptors and 15 groups of peaks. After determining the density of the receptors in each receptor-peak group we show in Table S7 only those with at least 0.5 density. All the groups of peaks that correspond to one group of receptors are aggregated to determine average frequency and average peak height.

Tissue-specific receptors evaluation

We have used the analysis pipeline, illustrated in Figure 6, separately on several tissue-specific studies (SRAs) from PanglaoDB annotated by DCS, from *Mus musculus* and *Homo sapiens* datasets. In order to compare the processed SRAs, we calculated the

Pearson correlation of the bootstrap frequencies, shown as a heatmap in Figure 3. The SRAs are ordered according to the Ward (Ward Jr, 1963) agglomerative clustering of the Euclidean similarity metric of the Pearson correlation of bootstrap frequencies.

The top tissue-specific receptors are shown as a heatmap in Figure 4, to show that some receptors can be found in all tissues but others are specific to one or a few tissues. SRA studies of the same tissue were combined by averaging the largest peak bootstrap frequencies. The lists of the top-10 receptors for each of the analyzed tissues and the bootstrap frequency cutoffs are given in Table S8.

Metascape enrichment analysis

For pathway enrichment analysis, Metascape (Zhou et al., 2019) calculates the p-value defined as the probability from the cumulative hypergeometric distribution and corrects it with the Bonferroni method to give the q-value.

ACKNOWLEDGEMENTS

This work was supported by National Institutes of Health (No. R01GM122085). We are very grateful for the statistical feedback received from Michael G. Walker and Karen Messer.

COMPETING INTERESTS

CP and GP own shares of Salgomed, Inc.

REFERENCES

Ackermann, M., S.E. Verleden, M. Kuehnel, A. Haverich, T. Welte, F. Laenger, A. Vanstapel, C. Werlein, H. Stark, A. Tzankov, W.W. Li, V.W. Li, S.J. Mentzer, and D. Jonigk. 2020. Pulmonary Vascular Endothelialitis, Thrombosis, and Angiogenesis in Covid-19. *The New England journal of medicine* 383:120-128.

Agre, P., and A.I. Leshner. 2010. Bridging science and society. In American Association for the Advancement of Science.

Almen, M.S., K.J. Nordstrom, R. Fredriksson, and H.B. Schioth. 2009. Mapping the human membrane proteome: a majority of the human membrane proteins can be classified according to function and evolutionary origin. *BMC biology* 7:50.

Antibody Society. 2020. Antibody therapeutics approved or in regulatory review in the EU or US. In.

Apte, R.S., D.S. Chen, and N. Ferrara. 2019. VEGF in Signaling and Disease: Beyond Discovery and Development. *Cell* 176:1248-1264.

Armingol, E., A. Officer, O. Harismendy, and N.E. Lewis. 2020. Deciphering cell–cell interactions and communication from gene expression. *Nature Reviews Genetics* 1-18.

Azoulay, P., J.S. Graff Zivin, and G. Manso. 2011. Incentives and creativity: evidence from the academic life sciences. *The RAND Journal of Economics* 42:527-554.

Begley, C.G., and L.M. Ellis. 2012. Drug development: Raise standards for preclinical cancer research. *Nature* 483:531-533.

Blanco-Melo, D., B.E. Nilsson-Payant, W.C. Liu, S. Uhl, D. Hoagland, R. Moller, T.X. Jordan, K. Oishi, M. Panis, D. Sachs, T.T. Wang, R.E. Schwartz, J.K. Lim, R.A. Albrecht, and B.R. tenOever. 2020. Imbalanced Host Response to SARS-CoV-2 Drives Development of COVID-19. *Cell* 181:1036-1045 e1039.

Boeckx, C., and K.K. Grohmann. 2013. The Cambridge handbook of biolinguistics. Cambridge University Press,

Cho, A. 2020. With to-do list checked off, U.S. physicists ask, 'What's next?'. *Science* 370:156-157.

Chung, A.S., and N. Ferrara. 2011. Developmental and pathological angiogenesis. *Annual review of cell and developmental biology* 27:563-584.

Dawkins, R. 1997. Climbing mount improbable. WW Norton & Company,

Domanskyi, S., A. Hakansson, T. Bertus, G. Paternostro, and C. Piermarocchi. 2021. Digital Cell Sorter (DCS): a cell type identification, anomaly detection, and Hopfield landscapes toolkit for single-cell transcriptomics. *PeerJ* 9:e10670.

Domanskyi, S., C. Piermarocchi, and G.I. Mias. 2019a. PyIOmica: longitudinal omics analysis and trend identification. *Bioinformatics* 36:2306-2307.

Domanskyi, S., A. Szedlak, N.T. Hawkins, J. Wang, G. Paternostro, and C. Piermarocchi. 2019b. Polled Digital Cell Sorter (p-DCS): Automatic identification of hematological cell types from single cell RNA-sequencing clusters. *BMC bioinformatics* 20:369.

Efron, B., and T. Hastie. 2016. Computer Age Statistical Inference. Cambridge University Press,

Feala, J.D., J. Cortes, P.M. Duxbury, A.D. McCulloch, C. Piermarocchi, and G. Paternostro. 2012. Statistical properties and robustness of biological controller-target networks. *PloS one* 7:e29374.

Feala, J.D., J. Cortes, P.M. Duxbury, C. Piermarocchi, A.D. McCulloch, and G. Paternostro. 2010. Systems approaches and algorithms for discovery of combinatorial therapies *WIREs Syst Biol Med* 181-193.

Felsenstein, J. 1985. Confidence limits on phylogenies: an approach using the bootstrap. *evolution* 39:783-791.

Ferrara, N., and A.P. Adamis. 2016. Ten years of anti-vascular endothelial growth factor therapy. *Nature reviews. Drug discovery* 15:385-403.

Ferrara, N., H.P. Gerber, and J. LeCouter. 2003. The biology of VEGF and its receptors. *Nature medicine* 9:669-676.

Ferrara, N., and R.S. Kerbel. 2005. Angiogenesis as a therapeutic target. *Nature* 438:967-974.

Franzén, O., and J.L.M. Björkegren. 2020. alona: a web server for single-cell RNA-seq analysis. *Bioinformatics* 36:3910-3912.

Franzén, O., L.-M. Gan, and J.L.M. Björkegren. 2019. PanglaoDB: a web server for exploration of mouse and human single-cell RNA sequencing data. *Database* 2019:

Fritsche, L.G., W. Igl, J.N.C. Bailey, F. Grassmann, S. Sengupta, J.L. Bragg-Gresham, K.P. Burdon, S.J. Hebbring, C. Wen, and M. Gorski. 2016. A large genome-wide association study of age-related macular degeneration highlights contributions of rare and common variants. *Nature genetics* 48:134-143.

Frye, M., M. Dierkes, V. Küppers, M. Vockel, J. Tomm, D. Zeuschner, J. Rossaint, A. Zarbock, G.Y. Koh, and K. Peters. 2015. Interfering with VE-PTP stabilizes endothelial junctions in vivo via Tie-2 in the absence of VE-cadherin. *Journal of Experimental Medicine* 212:2267-2287.

Giving Pledge. 2020. The Giving Pledge. In.

Gomes, T., and S.A. Teichmann. 2020. An antiviral response beyond immune cells. *Nature* 583:206-207.

Good, P. 2013. Permutation tests: a practical guide to resampling methods for testing hypotheses. Springer Science & Business Media,

Huang, J.K., D.E. Carlin, M.K. Yu, W. Zhang, J.F. Kreisberg, P. Tamayo, and T. Ideker. 2018. Systematic Evaluation of Molecular Networks for Discovery of Disease Genes. *Cell systems* 6:484-495 e485.

HuBmap, C. 2019. The human body at cellular resolution: the NIH Human Biomolecular Atlas Program. *Nature* 574:187-192.

Institute for Policy Studies. 2020. Gilded Giving 2020. In.

Jacob, F. 1977. Evolution and tinkering. *Science* 196:1161-1166.

Jacob, F., D. Perrin, C. Sanchez, and J. Monod. 1960. The Operon: a group of genes with the expression coordinated by an operator. *Comptes rendus hebdomadaires des seances de l'Academie des sciences* 250:1727-1729.

Kang, Y., A. Hodges, E. Ong, W. Roberts, C. Piermarocchi, and G. Paternostro. 2014. Identification of Drug Combinations Containing Imatinib for Treatment of BCR-ABL+ Leukemias. *PloS one* 9:e102221.

Karagiannis, P., K. Takahashi, M. Saito, Y. Yoshida, K. Okita, A. Watanabe, H. Inoue, J.K. Yamashita, M. Todani, M. Nakagawa, M. Osawa, Y. Yashiro, S. Yamanaka, and K. Osafune. 2019. Induced Pluripotent Stem Cells and Their Use in Human Models of Disease and Development. *Physiological reviews* 99:79-114.

Kaur, A., M.R. Webster, K. Marchbank, R. Behera, A. Ndoye, C.H. Kugel, V.M. Dang, J. Appleton, M.P. O'Connell, and P. Cheng. 2016. sFRP2 in the aged microenvironment drives melanoma metastasis and therapy resistance. *Nature* 532:250-254.

Köhler, R. 2012. Quantitative syntax analysis. Walter de Gruyter,

Komili, S., and P.A. Silver. 2008. Coupling and coordination in gene expression processes: a systems biology view. *Nature reviews. Genetics* 9:38-48.

Kong, D.H., M.R. Kim, J.H. Jang, H.J. Na, and S. Lee. 2017. A Review of Anti-Angiogenic Targets for Monoclonal Antibody Cancer Therapy. *International journal of molecular sciences* 18:

Krausgruber, T., N. Fortelny, V. Fife-Gernedl, M. Senekowitsch, L.C. Schuster, A. Lercher, A. Nemc, C. Schmidl, A.F. Rendeiro, A. Bergthaler, and C. Bock. 2020. Structural cells are key regulators of organ-specific immune responses. *Nature* 583:296-302.

Kuenzi, B.M., J. Park, S.H. Fong, K.S. Sanchez, J. Lee, J.F. Kreisberg, J. Ma, and T. Ideker. 2020. Predicting drug response and synergy using a deep learning model of human Cancer cells. *Cancer cell* 38:672-684. e676.

Latchman, D. 2015. Gene Control, Second Edition. CRC Press LLC,

Li, Q., P. Li, N. Biswas, L. Liu, H. Xin, G. Paternostro, C. Piermarocchi, S. Domanskyi, R.K. Wang, and N. Ferrara. 2021. Identification of LIF as a mitogen for choroidal endothelial cells: cell-type specific STAT3 signaling determines whether LIF stimulates or inhibits endothelial cell growth. *submitted for publication*

Manly, B.F. 2006. Randomization, bootstrap and Monte Carlo methods in biology. CRC press,

Maxmen, A. 2020. Untapped potential: More US labs could be providing tests for coronavirus. In Nature.

Morens, D.M., and A.S. Fauci. 2020. Emerging Pandemic Diseases: How We Got to COVID-19. *Cell* 182:1077-1092.

Murray, F., P. Aghion, M. Dewatripont, J. Kolev, and S. Stern. 2016. Of mice and academics: Examining the effect of openness on innovation. *American Economic Journal: Economic Policy* 8:212-252.

National Academies of Sciences, E., and Medicine. 2018. Open Science by Design: Realizing a Vision for 21st Century Research. National Academies Press (US), Washington (DC).

Nowak, M.A., M.C. Boerlijst, J. Cooke, and J.M. Smith. 1997. Evolution of genetic redundancy. *Nature* 388:167-171.

Nurse, P. 2020. What Is Life?: understand biology in five steps. Scribe Publications Pty Limited,

O'Clery, C. 2013. The Billionaire Who Wasn't: How Chuck Feeney Secretly Made and Gave Away a Fortune. PublicAffairs,

Ohno, S. 2013. Evolution by gene duplication. Springer Science & Business Media,

Ong, E., A. Szedlak, Y. Kang, P. Smith, N. Smith, M. McBride, D. Finlay, K. Vuori, J. Mason, E.D. Ball, C. Piermarocchi, and G. Paternostro. 2015. A scalable method for molecular network reconstruction identifies properties of targets and mutations in acute myeloid leukemia. *Journal of computational biology : a journal of computational molecular cell biology* 22:266-288.

Palca, J. 1992. The genome project: life after Watson. *Science* 256:956-958.

Potente, M., H. Gerhardt, and P. Carmeliet. 2011. Basic and therapeutic aspects of angiogenesis. *Cell* 146:873-887.

Prinz, F., T. Schlange, and K. Asadullah. 2011. Believe it or not: how much can we rely on published data on potential drug targets? *Nature reviews. Drug discovery* 10:712.

Regev, A., S.A. Teichmann, E.S. Lander, I. Amit, C. Benoist, E. Birney, B. Bodenmiller, P. Campbell, P. Carninci, M. Clatworthy, H. Clevers, B. Deplancke, I. Dunham, J. Eberwine, R. Eils, W. Enard, A. Farmer, L. Fugger, B. Gottgens, N. Hacohen, M. Haniffa, M. Hemberg, S. Kim, P. Klenerman, A. Kriegstein, E. Lein, S. Linnarsson, E. Lundberg, J. Lundeberg, P. Majumder, J.C. Marioni, M. Merad, M. Mhlanga, M. Nawijn, M. Netea, G. Nolan, D. Pe'er, A. Phillipakis, C.P. Ponting, S. Quake, W. Reik, O. Rozenblatt-Rosen, J. Sanes, R. Satija, T.N. Schumacher, A. Shalek, E. Shapiro, P. Sharma, J.W. Shin, O. Stegle, M. Stratton, M.J.T. Stubbington, F.J. Theis, M. Uhlen, A. van Oudenaarden, A. Wagner, F. Watt, J. Weissman, B. Wold, R. Xavier, N. Yosef, and P. Human Cell Atlas Meeting. 2017. The Human Cell Atlas. *eLife* 6:

Rieckmann, J.C., R. Geiger, D. Hornburg, T. Wolf, K. Kveler, D. Jarrossay, F. Sallusto, S.S. Shen-Orr, A. Lanzavecchia, M. Mann, and F. Meissner. 2017. Social network architecture of human immune cells unveiled by quantitative proteomics. *Nature immunology*

Sakurai, A., J. Gavard, Y. Annas-Linhares, J.R. Basile, P. Amornphimoltham, T.R. Palmby, H. Yagi, F. Zhang, P.A. Randazzo, X. Li, R. Weigert, and J.S. Gutkind. 2010. Semaphorin 3E initiates antiangiogenic signaling through plexin D1 by regulating Arf6 and R-Ras. *Molecular and cellular biology* 30:3086-3098.

Scannell, J.W., A. Blanckley, H. Boldon, and B. Warrington. 2012. Diagnosing the decline in pharmaceutical R&D efficiency. *Nature reviews. Drug discovery* 11:191-200.

Silva, J.S., and S. Tenreyro. 2006. The log of gravity. *The Review of Economics and statistics* 88:641-658.

Singh, S. 1999. The code book. Doubleday, New York. xiii, 402 p. pp.

Soltis, P.S., and D.E. Soltis. 2003. Applying the bootstrap in phylogeny reconstruction. *Statistical Science* 256-267.

Strohl, W.R. 2018. Current progress in innovative engineered antibodies. *Protein & cell* 9:86-120.

Sturtzel, C., J. Testori, B. Schweighofer, M. Bilban, and E. Hofer. 2014. The transcription factor MEF2C negatively controls angiogenic sprouting of endothelial cells depending on oxygen. *PloS one* 9:e101521.

Szedlak, A., N. Smith, L. Liu, G. Paternostro, and C. Piermarocchi. 2016. Evolutionary and Topological Properties of Genes and Community Structures in Human Gene Regulatory Networks. *PLoS computational biology* 12:e1005009.

The RECOVERY Collaborative Group. 2021. Dexamethasone in Hospitalized Patients with Covid-19. *New Engl J Med* 384:693-704.

Tzima, E., M. Irani-Tehrani, W.B. Kiosses, E. Dejana, D.A. Schultz, B. Engelhardt, G. Cao, H. DeLisser, and M.A. Schwartz. 2005. A mechanosensory complex that mediates the endothelial cell response to fluid shear stress. *Nature* 437:426-431.

Van Noorden, R., B. Maher, and R. Nuzzo. 2014. The top 100 papers. *Nature News* 514:550.

Voigt, A.P., K. Mulfaul, N.K. Mullin, M.J. Flamme-Wiese, J.C. Giacalone, E.M. Stone, B.A. Tucker, T.E. Scheetz, and R.F. Mullins. 2019. Single-cell transcriptomics of the human retinal pigment epithelium and choroid in health and macular degeneration. *Proc Natl Acad Sci U S A* 116:24100-24107.

Wagner, A. 2005. Robustness and evolvability in living systems. Princeton University Presss, Princeton.

Ward Jr, J.H. 1963. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association* 58:236-244.

Williams, H.L. 2013. Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy* 121:1-27.

Wooldridge, J.M. 1999. Quasi-likelihood methods for count data. *Handbook of applied econometrics* 2:352-406.

Yancopoulos, G.D., S. Davis, N.W. Gale, J.S. Rudge, S.J. Wiegand, and J. Holash. 2000. Vascular-specific growth factors and blood vessel formation. *Nature* 407:242-248.

Yang, W.J., J. Hu, A. Uemura, F. Tetzlaff, H.G. Augustin, and A. Fischer. 2015. Semaphorin-3C signals through Neuropilin-1 and PlexinD1 receptors to inhibit pathological angiogenesis. *EMBO Mol Med* 7:1267-1284.

Zhou, Y., B. Zhou, L. Pache, M. Chang, A.H. Khodabakhshi, O. Tanaseichuk, C. Benner, and S.K. Chanda. 2019. Metascape provides a biologist-oriented resource for the analysis of systems-level datasets. *Nat Commun* 10:1523.